

Small cryptographic bytecode

D. J. Bernstein

elaborating on an idea from  
Adam Langley

“Line search”:

trying to find minimum of  
function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find  
minimum in interval  $[x_0, x_1]$ :

Replace interval with either  
 $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;  
try to make sensible choice.

Iterate many times.

Small cryptographic bytecode

D. J. Bernstein

elaborating on an idea from  
Adam Langley

1

“Line search”:

trying to find minimum of  
function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find  
minimum in interval  $[x_0, x_1]$ :

Replace interval with either  
 $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;  
try to make sensible choice.

Iterate many times.

Can try to reduce #iterations  
using smarter models of  $f$ :  
see, e.g., “secant method”.

2

Small cryptographic bytecode

D. J. Bernstein

elaborating on an idea from  
Adam Langley

1

“Line search”:

trying to find minimum of  
function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find  
minimum in interval  $[x_0, x_1]$ :

Replace interval with either  
 $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;  
try to make sensible choice.

Iterate many times.

Can try to reduce #iterations  
using smarter models of  $f$ :  
see, e.g., “secant method”.

Harder when  $f$  varies more.

2

cryptographic bytecode

ernstein

ing on an idea from

angley

1

“Line search” :

trying to find minimum of  
function  $f$  defined on  $x$ -line.

e.g. “Bisection” , trying to find  
minimum in interval  $[x_0, x_1]$ :

Replace interval with either  
 $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;  
try to make sensible choice.

Iterate many times.

Can try to reduce #iterations  
using smarter models of  $f$ :  
see, e.g., “secant method” .

Harder when  $f$  varies more.

2

How to  
 $f$  defined

“Gradient  
Starting  
try to fig  
where  $f$

ic bytecode

idea from

1

“Line search”:

trying to find minimum of  
function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find  
minimum in interval  $[x_0, x_1]$ :

Replace interval with either  
 $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;  
try to make sensible choice.

Iterate many times.

Can try to reduce #iterations  
using smarter models of  $f$ :  
see, e.g., “secant method”.

Harder when  $f$  varies more.

2

How to find minimum  
 $f$  defined on  $(x, y)$

“Gradient descent”  
Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases

1

“Line search” :

trying to find minimum of function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find minimum in interval  $[x_0, x_1]$ :

Replace interval with either  $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ; try to make sensible choice.

Iterate many times.

Can try to reduce #iterations using smarter models of  $f$ : see, e.g., “secant method” .

Harder when  $f$  varies more.

2

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent” :

Starting from  $(x_0, y_0)$ , try to figure out direction where  $f$  decreases fastest.

“Line search”:

trying to find minimum of function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find minimum in interval  $[x_0, x_1]$ :

Replace interval with either  $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ; try to make sensible choice.

Iterate many times.

Can try to reduce #iterations using smarter models of  $f$ : see, e.g., “secant method”.

Harder when  $f$  varies more.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ , try to figure out direction where  $f$  decreases fastest.

“Line search”:

trying to find minimum of function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find minimum in interval  $[x_0, x_1]$ :

Replace interval with either  $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;

try to make sensible choice.

Iterate many times.

Can try to reduce #iterations

using smarter models of  $f$ :

see, e.g., “secant method”.

Harder when  $f$  varies more.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,

try to figure out direction

where  $f$  decreases fastest.

Could do line search to find minimum in that direction.

Then find a new direction.

“Line search”:

trying to find minimum of function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find minimum in interval  $[x_0, x_1]$ :

Replace interval with either  $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ;

try to make sensible choice.

Iterate many times.

Can try to reduce #iterations

using smarter models of  $f$ :

see, e.g., “secant method”.

Harder when  $f$  varies more.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,

try to figure out direction

where  $f$  decreases fastest.

Could do line search to find minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

“Line search”:

trying to find minimum of function  $f$  defined on  $x$ -line.

e.g. “Bisection”, trying to find minimum in interval  $[x_0, x_1]$ :

Replace interval with either  $[x_0, (x_0 + x_1)/2]$  or  $[(x_0 + x_1)/2, x_1]$ ; try to make sensible choice.

Iterate many times.

Can try to reduce #iterations using smarter models of  $f$ : see, e.g., “secant method”.

Harder when  $f$  varies more.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ , try to figure out direction where  $f$  decreases fastest.

Could do line search to find minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction; line search in  $y$  direction; repeat.

Search” :

to find minimum of  
 $f$  defined on  $x$ -line.

“bisection”, trying to find  
min in interval  $[x_0, x_1]$ :

interval with either  
 $[(x_0 - x_1)/2, x_0]$  or  $[(x_0 + x_1)/2, x_1]$ ;

make sensible choice.

many times.

to reduce #iterations

quadratic models of  $f$ :

, “secant method” .

when  $f$  varies more.

2

How to find minimum of function  
 $f$  defined on  $(x, y)$ -plane?

“Gradient descent” :

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

3

Keccak

Goal: Fa  
on a Co

You star  
impleme

2

mum of  
on  $x$ -line.  
trying to find  
val  $[x_0, x_1]$ :  
with either  
 $[(x_0 + x_1)/2, x_1]$ ;  
le choice.  
s.  
#iterations  
dels of  $f$ :  
method".  
ries more.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

3

Keccak optimization

Goal: Fastest C code  
on a Cortex-M4 C

You start with sim  
implementing Kecc

2

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

3

## Keccak optimization

Goal: Fastest C code for Keccak  
on a Cortex-M4 CPU core.

You start with simple C code  
implementing Keccak.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

## Keccak optimization

Goal: Fastest C code for Keccak  
on a Cortex-M4 CPU core.

You start with simple C code  
implementing Keccak.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

## Keccak optimization

Goal: Fastest C code for Keccak  
on a Cortex-M4 CPU core.

You start with simple C code  
implementing Keccak.

You compile it; see how fast it is;  
modify it to try to make it faster;  
repeat; eventually stop trying.

How to find minimum of function  $f$  defined on  $(x, y)$ -plane?

“Gradient descent”:

Starting from  $(x_0, y_0)$ ,  
try to figure out direction  
where  $f$  decreases fastest.

Could do line search to find  
minimum in that direction.

Then find a new direction.

Better: Step down that direction.

Then find a new direction.

Silly: Line search in  $x$  direction;  
line search in  $y$  direction; repeat.

## Keccak optimization

Goal: Fastest C code for Keccak  
on a Cortex-M4 CPU core.

You start with simple C code  
implementing Keccak.

You compile it; see how fast it is;  
modify it to try to make it faster;  
repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it,  
and care about its speed.

find minimum of function  
d on  $(x, y)$ -plane?

“steepest descent”:

start from  $(x_0, y_0)$ ,  
figure out direction  
in which function  
decreases fastest.

do a line search to find  
minimum in that direction.  
then find a new direction.

Step down that direction.  
find a new direction.

do a line search in  $x$  direction;  
do a line search in  $y$  direction; repeat.

3

## Keccak optimization

Goal: Fastest C code for Keccak  
on a Cortex-M4 CPU core.

You start with simple C code  
for implementing Keccak.

You compile it; see how fast it is;  
modify it to try to make it faster;  
repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it,  
and care about its speed.

4

Compile  
your Keccak

3

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

4

Compiler writer leads you to your Keccak Cortex-M4 code.

3

nction

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

ction.

tion;

peat.

4

Compiler writer learns about your Keccak Cortex-M4 C code

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

Compiler writer learns about your Keccak Cortex-M4 C code.

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new compiler. Whole process repeats.

## Keccak optimization

Goal: Fastest C code for Keccak on a Cortex-M4 CPU core.

You start with simple C code implementing Keccak.

You compile it; see how fast it is; modify it to try to make it faster; repeat; eventually stop trying.

You publish your fastest code.

Maybe lots of people use it, and care about its speed.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

## optimization

fastest C code for Keccak  
Cortex-M4 CPU core.

Start with simple C code  
implementing Keccak.

Compile it; see how fast it is;  
try to try to make it faster;  
eventually stop trying.

Publish your fastest code.

Lots of people use it,  
care about its speed.

4

Compiler writer learns about  
your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to  
make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new  
compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

5

Define *f*  
code  $x$  v

4

on  
ode for Keccak  
PU core.  
ple C code  
cak.  
e how fast it is;  
o make it faster;  
stop trying.  
fastest code.  
ple use it,  
speed.

Compiler writer learns about  
your Keccak Cortex-M4 C code.  
Compiles it; sees how fast it is.  
Modifies *compiler* to try to  
make the compiled code faster.  
Repeats; eventually stops trying.  
Publishes a new compiler version.  
Later: Maybe you try the new  
compiler. Whole process repeats.  
You treat compiler as constant.  
Compiler treats code as constant.

5

Define  $f(x, y)$  as t  
code  $x$  with comp

4

Compiler writer learns about  
your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to  
make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new  
compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

5

Define  $f(x, y)$  as time taken  
code  $x$  with compiler  $y$ .

Compiler writer learns about  
your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.  
Modifies *compiler* to try to  
make the compiled code faster.  
Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new  
compiler. Whole process repeats.

You treat compiler as constant.  
Compiler treats code as constant.

Define  $f(x, y)$  as time taken by  
code  $x$  with compiler  $y$ .

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer learns about  
 your Keccak Cortex-M4 C code.  
 Compiles it; sees how fast it is.  
 Modifies *compiler* to try to  
 make the compiled code faster.  
 Repeats; eventually stops trying.  
 Publishes a new compiler version.  
 Later: Maybe you try the new  
 compiler. Whole process repeats.  
 You treat compiler as constant.  
 Compiler treats code as constant.

Define  $f(x, y)$  as time taken by  
 code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this  
 line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this  
 line search in  $y$  direction.

Compiler writer learns about your Keccak Cortex-M4 C code.

Compiles it; sees how fast it is.

Modifies *compiler* to try to make the compiled code faster.

Repeats; eventually stops trying.

Publishes a new compiler version.

Later: Maybe you try the new compiler. Whole process repeats.

You treat compiler as constant.

Compiler treats code as constant.

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

Compiler writer learns about  
gccak Cortex-M4 C code.  
Runs it; sees how fast it is.  
Tries *compiler* to try to  
get the compiled code faster.  
Eventually stops trying.  
Finds a new compiler version.  
Maybe you try the new  
one. Whole process repeats.  
Compiler treats compiler as constant.  
Compiler treats code as constant.

5

Define  $f(x, y)$  as time taken by  
code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this  
line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this  
line search in  $y$  direction.

This whole approach is silly.

6

$\min\{f(x, y)\}$   
fastest k

5

learns about  
 ex-M4 C code.  
 how fast it is.  
 to try to  
 d code faster.  
 y stops trying.  
 ompiler version.  
 try the new  
 process repeats.  
 r as constant.  
 ode as constant.

Define  $f(x, y)$  as time taken by  
 code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this  
 line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this  
 line search in  $y$  direction.

This whole approach is silly.

6

$\min\{f(x, y)\}$  is th  
 fastest Keccak Co

5

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

6

$\min\{f(x, y)\}$  is the time taken by the fastest Keccak Cortex-M4 as

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Define  $f(x, y)$  as time taken by code  $x$  with compiler  $y$ .

$x_0$ : initial code.

$y_0$ : initial compiler.

You try to minimize  $f(x, y_0)$ .

$x_1$ : new code from this line search in  $x$  direction.

Compiler writer:  $f(x_1, y)$ .

$y_1$ : new compiler from this line search in  $y$  direction.

This whole approach is silly.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say "C code"?

End user doesn't need C.

$f(x, y)$  as time taken by  
with compiler  $y$ .

al code.

al compiler.

to minimize  $f(x, y_0)$ .

code from this

ch in  $x$  direction.

r writer:  $f(x_1, y)$ .

compiler from this

ch in  $y$  direction.

ole approach is silly.

6

$\min\{f(x, y)\}$  is the time taken by  
fastest Keccak Cortex-M4 asm.

Slowly bouncing between  
 $x$ -line searches,  $y$ -line searches is  
a silly way to approach this min.

Clearly min can be achieved by  
many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other  
languages: which language  
makes min easiest to find?

Why did goal say "C code"?

End user doesn't need C.

7

Does en

time taken by  
filer  $y$ .

r.

ize  $f(x, y_0)$ .

n this  
rection.

$f(x_1, y)$ .

from this  
rection.

ch is silly.

6

$\min\{f(x, y)\}$  is the time taken by  
fastest Keccak Cortex-M4 asm.

Slowly bouncing between  
 $x$ -line searches,  $y$ -line searches is  
a silly way to approach this min.

Clearly min can be achieved by  
many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other  
languages: which language  
makes min easiest to find?

Why did goal say "C code"?

End user doesn't need C.

7

Does end user need

6

by

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say "C code"?

End user doesn't need C.

7

Does end user need Cortex-M4

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say “C code”?

End user doesn't need C.

Does end user need Cortex-M4?

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say “C code”?

End user doesn't need C.

Does end user need Cortex-M4?  
CPU designer learns about your Keccak Cortex-M4 asm.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say “C code”?

End user doesn't need C.

Does end user need Cortex-M4?  
CPU designer learns about your Keccak Cortex-M4 asm.

Modifies the CPU design to try to make this code faster.

Repeats; eventually stops trying.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say “C code”?

End user doesn't need C.

Does end user need Cortex-M4?  
CPU designer learns about your Keccak Cortex-M4 asm.

Modifies the CPU design to try to make this code faster.  
Repeats; eventually stops trying.

Years later, sells a new CPU.  
You reoptimize for this CPU.

$\min\{f(x, y)\}$  is the time taken by fastest Keccak Cortex-M4 asm.

Slowly bouncing between  $x$ -line searches,  $y$ -line searches is a silly way to approach this min.

Clearly min can be achieved by many different pairs  $(x, y)$ .

Which pair is easiest to find?

Generalize from C to other languages: which language makes min easiest to find?

Why did goal say "C code"?

End user doesn't need C.

Does end user need Cortex-M4?  
CPU designer learns about your Keccak Cortex-M4 asm.

Modifies the CPU design to try to make this code faster.  
Repeats; eventually stops trying.

Years later, sells a new CPU.  
You reoptimize for this CPU.

*Sometimes* CPUs try extending or replacing instruction set, but this is poorly coordinated with programmers, compiler writers.

$(x, y)$  is the time taken by  
Keccak Cortex-M4 asm.

bouncing between

searches,  $y$ -line searches is

way to approach this min.

min can be achieved by

different pairs  $(x, y)$ .

pair is easiest to find?

ize from C to other

es: which language

min easiest to find?

goal say "C code"?

r doesn't need C.

7

Does end user need Cortex-M4?  
CPU designer learns about your  
Keccak Cortex-M4 asm.

Modifies the CPU design to  
try to make this code faster.  
Repeats; eventually stops trying.

Years later, sells a new CPU.

You reoptimize for this CPU.

*Sometimes* CPUs try extending  
or replacing instruction set, but  
this is poorly coordinated with  
programmers, compiler writers.

8

Generaliz  
 $f(x, y)$  i  
code  $x$  o

If compi  
asm  $y(x)$   
 $f(x, y) =$

7

the time taken by  
Cortex-M4 asm.

between

line searches is  
to reach this min.

is achieved by

searchers  $(x, y)$ .

How best to find?

How to other

language

to find?

“C code”?

Do we need C.

Does end user need Cortex-M4?

CPU designer learns about your

Keccak Cortex-M4 asm.

Modifies the CPU design to

try to make this code faster.

Repeats; eventually stops trying.

Years later, sells a new CPU.

You reoptimize for this CPU.

*Sometimes* CPUs try extending

or replacing instruction set, but

this is poorly coordinated with

programmers, compiler writers.

8

Generalize  $f(x, y)$

$f(x, y)$  is time taken

to run code  $x$  on platform

If compiler  $y$  on code

asm  $y(x)$  for Cortex

$f(x, y) = f(y(x), \dots)$

7

Does end user need Cortex-M4?  
 CPU designer learns about your  
 Keccak Cortex-M4 asm.

Modifies the CPU design to  
 try to make this code faster.  
 Repeats; eventually stops trying.

Years later, sells a new CPU.  
 You reoptimize for this CPU.

*Sometimes* CPUs try extending  
 or replacing instruction set, but  
 this is poorly coordinated with  
 programmers, compiler writers.

8

Generalize  $f(x, y)$  definition  
 $f(x, y)$  is time taken by  
 code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
 asm  $y(x)$  for Cortex-M4:  
 $f(x, y) = f(y(x), \text{Cortex-M4})$

Does end user need Cortex-M4?

CPU designer learns about your Keccak Cortex-M4 asm.

Modifies the CPU design to try to make this code faster.  
Repeats; eventually stops trying.

Years later, sells a new CPU.

You reoptimize for this CPU.

*Sometimes* CPUs try extending or replacing instruction set, but this is poorly coordinated with programmers, compiler writers.

Generalize  $f(x, y)$  definition:

$f(x, y)$  is time taken by code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces asm  $y(x)$  for Cortex-M4:

$f(x, y) = f(y(x), \text{Cortex-M4})$ .

Does end user need Cortex-M4?  
 CPU designer learns about your  
 Keccak Cortex-M4 asm.

Modifies the CPU design to  
 try to make this code faster.  
 Repeats; eventually stops trying.

Years later, sells a new CPU.  
 You reoptimize for this CPU.

*Sometimes* CPUs try extending  
 or replacing instruction set, but  
 this is poorly coordinated with  
 programmers, compiler writers.

Generalize  $f(x, y)$  definition:  
 $f(x, y)$  is time taken by  
 code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
 asm  $y(x)$  for Cortex-M4:  
 $f(x, y) = f(y(x), \text{Cortex-M4})$ .

Without the CPU changing:  
 Minimize  $f(a, \text{Cortex-M4})$ .  
 Search for  $(x, y)$  with  $y(x) = a$ .

Does end user need Cortex-M4?

CPU designer learns about your Keccak Cortex-M4 asm.

Modifies the CPU design to try to make this code faster.  
Repeats; eventually stops trying.

Years later, sells a new CPU.  
You reoptimize for this CPU.

*Sometimes* CPUs try extending or replacing instruction set, but this is poorly coordinated with programmers, compiler writers.

Generalize  $f(x, y)$  definition:

$f(x, y)$  is time taken by code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces asm  $y(x)$  for Cortex-M4:

$$f(x, y) = f(y(x), \text{Cortex-M4}).$$

Without the CPU changing:

Minimize  $f(a, \text{Cortex-M4})$ .

Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:

View  $a$  as a constant;

try to minimize  $f(a, y)$ .

Silly optimization approach.

...d user need Cortex-M4?  
...signer learns about your  
...Cortex-M4 asm.

...s the CPU design to  
...ake this code faster.  
...; eventually stops trying.

...ter, sells a new CPU.  
...optimize for this CPU.

...nes CPUs try extending  
...cing instruction set, but  
...poorly coordinated with  
...mers, compiler writers.

8

Generalize  $f(x, y)$  definition:  
 $f(x, y)$  is time taken by  
code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
asm  $y(x)$  for Cortex-M4:  
 $f(x, y) = f(y(x), \text{Cortex-M4})$ .

Without the CPU changing:  
Minimize  $f(a, \text{Cortex-M4})$ .  
Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:  
View  $a$  as a constant;  
try to minimize  $f(a, y)$ .  
Silly optimization approach.

9

“I know  
I’ve deve  
that con  
This circ

and Cortex-M4?

ns about your  
t asm.

design to  
ode faster.  
y stops trying.

new CPU.  
r this CPU.

try extending  
ction set, but  
dinated with  
mpiler writers.

8

Generalize  $f(x, y)$  definition:

$f(x, y)$  is time taken by  
code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
asm  $y(x)$  for Cortex-M4:

$f(x, y) = f(y(x), \text{Cortex-M4})$ .

Without the CPU changing:

Minimize  $f(a, \text{Cortex-M4})$ .

Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:

View  $a$  as a constant;

try to minimize  $f(a, y)$ .

Silly optimization approach.

9

“I know the minim  
I’ve developed the  
that computes Ke  
This circuit is my

Generalize  $f(x, y)$  definition:

$f(x, y)$  is time taken by  
code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
asm  $y(x)$  for Cortex-M4:

$$f(x, y) = f(y(x), \text{Cortex-M4}).$$

Without the CPU changing:

Minimize  $f(a, \text{Cortex-M4})$ .

Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:

View  $a$  as a constant;

try to minimize  $f(a, y)$ .

Silly optimization approach.

“I know the minimum!

I’ve developed the fastest circuit  
that computes Keccak.

This circuit is my CPU.”

Generalize  $f(x, y)$  definition:  
 $f(x, y)$  is time taken by  
code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
asm  $y(x)$  for Cortex-M4:  
 $f(x, y) = f(y(x), \text{Cortex-M4})$ .

Without the CPU changing:  
Minimize  $f(a, \text{Cortex-M4})$ .  
Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:  
View  $a$  as a constant;  
try to minimize  $f(a, y)$ .  
Silly optimization approach.

“I know the minimum!  
I’ve developed the fastest circuit  
that computes Keccak.  
This circuit is my CPU.”

Generalize  $f(x, y)$  definition:  
 $f(x, y)$  is time taken by  
 code  $x$  on platform  $y$ .

If compiler  $y$  on code  $x$  produces  
 asm  $y(x)$  for Cortex-M4:  
 $f(x, y) = f(y(x), \text{Cortex-M4})$ .

Without the CPU changing:  
 Minimize  $f(a, \text{Cortex-M4})$ .  
 Search for  $(x, y)$  with  $y(x) = a$ .

Typical CPU designer:  
 View  $a$  as a constant;  
 try to minimize  $f(a, y)$ .  
 Silly optimization approach.

“I know the minimum!  
 I’ve developed the fastest circuit  
 that computes Keccak.  
 This circuit is my CPU.”

Wait a minute: “CPU” concept  
 is more restrictive than “chip”.

Perspective of CPU designer:  
 This chip can do anything!

People want this chip to support  
 SHA-1, SHA-2, SHA-3, SHAMir;  
 all sorts of block ciphers;  
 public-key cryptosystems;  
 non-cryptographic computations.

minimize  $f(x, y)$  definition:

minimize time taken by

on platform  $y$ .

minimize time taken by

for Cortex-M4:

$= f(y(x), \text{Cortex-M4})$ .

minimize time taken by

for Cortex-M4.

minimize time taken by for  $(x, y)$  with  $y(x) = a$ .

minimize time taken by

as a constant;

minimize  $f(a, y)$ .

minimization approach.

“I know the minimum!

I’ve developed the fastest circuit that computes Keccak.

This circuit is my CPU.”

Wait a minute: “CPU” concept is more restrictive than “chip”.

Perspective of CPU designer:

This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding time taken by

(“Keccak”)

adds are

Adding time taken by

for desired

adds even

definition:

en by

n  $y$ .

ode  $x$  produces

ex-M4:

Cortex-M4).

changing:

tex-M4).

with  $y(x) = a$ .

gner:

ant;

$a, y$ ).

approach.

“I know the minimum!

I’ve developed the fastest circuit  
that computes Keccak.

This circuit is my CPU.”

Wait a minute: “CPU” concept  
is more restrictive than “chip”.

Perspective of CPU designer:

This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding fast Keccak

(“Keccak coprocessor”  
adds area to CPU.

Adding fast coprocessor

for desired mix of

adds even more area

“I know the minimum!  
I’ve developed the fastest circuit  
that computes Keccak.  
This circuit is my CPU.”

Wait a minute: “CPU” concept  
is more restrictive than “chip” .

Perspective of CPU designer:  
This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding fast Keccak circuit  
( “Keccak coprocessor” ) to C  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU

“I know the minimum!  
I’ve developed the fastest circuit  
that computes Keccak.  
This circuit is my CPU.”

Wait a minute: “CPU” concept  
is more restrictive than “chip” .

Perspective of CPU designer:  
This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding fast Keccak circuit  
( “Keccak coprocessor” ) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

“I know the minimum!  
I’ve developed the fastest circuit  
that computes Keccak.  
This circuit is my CPU.”

Wait a minute: “CPU” concept  
is more restrictive than “chip” .

Perspective of CPU designer:  
This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding fast Keccak circuit  
( “Keccak coprocessor” ) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

“I know the minimum!  
I’ve developed the fastest circuit  
that computes Keccak.  
This circuit is my CPU.”

Wait a minute: “CPU” concept  
is more restrictive than “chip” .

Perspective of CPU designer:  
This chip can do anything!

People want this chip to support  
SHA-1, SHA-2, SHA-3, SHAMir;  
all sorts of block ciphers;  
public-key cryptosystems;  
non-cryptographic computations.

Adding fast Keccak circuit  
( “Keccak coprocessor” ) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn’t reflect general case.

the minimum!  
 developed the fastest circuit  
 computes Keccak.  
 "This circuit is my CPU."  
 minute: "CPU" concept  
 restrictive than "chip".  
 perspective of CPU designer:  
 chip can do anything!  
 want this chip to support  
 SHA-2, SHA-3, SHAMir;  
 of block ciphers;  
 key cryptosystems;  
 cryptographic computations.

Adding fast Keccak circuit  
 ("Keccak coprocessor") to CPU  
 adds area to CPU.

Adding fast coprocessors  
 for desired mix of operations  
 adds even more area to CPU.

For same CPU area,  
 obtain much better throughput  
 by building many copies  
 of original CPU core  
 without these coprocessors.

Fast Keccak chip is special case.  
 Doesn't reflect general case.

CPU des  
 What is  
 for a spe  
 within a

num!  
 fastest circuit  
 Keccak.  
 CPU.”  
 CPU” concept  
 than “chip” .  
 CPU designer:  
 anything!  
 chip to support  
 SHA-3, SHAmir;  
 ciphers;  
 systems;  
 computations.

Adding fast Keccak circuit  
 (“Keccak coprocessor”) to CPU  
 adds area to CPU.

Adding fast coprocessors  
 for desired mix of operations  
 adds even more area to CPU.

For same CPU area,  
 obtain much better throughput  
 by building many copies  
 of original CPU core  
 without these coprocessors.

Fast Keccak chip is special case.  
 Doesn't reflect general case.

CPU designer's mo  
 What is best perfor  
 for a specified mix  
 within a particular

Adding fast Keccak circuit  
(“Keccak coprocessor”) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn't reflect general case.

CPU designer's metric:  
What is best performance  
for a specified mix of operations  
within a particular CPU area

Adding fast Keccak circuit  
(“Keccak coprocessor”) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn't reflect general case.

CPU designer's metric:  
What is best performance  
for a specified mix of operations  
within a particular CPU area?

Adding fast Keccak circuit  
(“Keccak coprocessor”) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn't reflect general case.

CPU designer's metric:  
What is best performance  
for a specified mix of operations  
within a particular CPU area?

CPU designer is much more likely  
to consider incorporating a  
**small** Keccak coprocessor.

Adding fast Keccak circuit  
(“Keccak coprocessor”) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn't reflect general case.

CPU designer's metric:  
What is best performance  
for a specified mix of operations  
within a particular CPU area?

CPU designer is much more likely  
to consider incorporating a  
**small** Keccak coprocessor.

“So we should design the  
*smallest* Keccak circuit?”

Adding fast Keccak circuit  
(“Keccak coprocessor”) to CPU  
adds area to CPU.

Adding fast coprocessors  
for desired mix of operations  
adds even more area to CPU.

For same CPU area,  
obtain much better throughput  
by building many copies  
of original CPU core  
without these coprocessors.

Fast Keccak chip is special case.  
Doesn't reflect general case.

CPU designer's metric:  
What is best performance  
for a specified mix of operations  
within a particular CPU area?

CPU designer is much more likely  
to consider incorporating a  
**small** Keccak coprocessor.

“So we should design the  
*smallest* Keccak circuit?”

—Maybe, but will this extreme  
be faster than using existing CPU  
instructions without coprocessor?

fast Keccak circuit  
 k coprocessor”) to CPU  
 a to CPU.

fast coprocessors  
 ed mix of operations  
 en more area to CPU.

e CPU area,  
 much better throughput  
 ing many copies  
 al CPU core  
 these coprocessors.

ccak chip is special case.  
 reflect general case.

CPU designer’s metric:  
 What is best performance  
 for a specified mix of operations  
 within a particular CPU area?

CPU designer is much more likely  
 to consider incorporating a  
**small** Keccak coprocessor.

“So we should design the  
*smallest* Keccak circuit?”

—Maybe, but will this extreme  
 be faster than using existing CPU  
 instructions without coprocessor?

Intel typ  
 quite lar  
 32KB L1  
 32KB L1  
 several f  
 many di  
 out-of-o

“So it’s  
 to add in  
 for my f

CPU designer's metric:

What is best performance for a specified mix of operations within a particular CPU area?

CPU designer is much more likely to consider incorporating a **small** Keccak coprocessor.

“So we should design the *smallest* Keccak circuit?”

—Maybe, but will this extreme be faster than using existing CPU instructions without coprocessor?

Intel typically designs quite large CPU cores with 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, and an out-of-order unit,

“So it's small cost to add instruction-  
for my favorite cry

CPU

s

J.

put

case.

CPU designer's metric:

What is best performance for a specified mix of operations within a particular CPU area?

CPU designer is much more likely to consider incorporating a **small** Keccak coprocessor.

“So we should design the *smallest* Keccak circuit?”

—Maybe, but will this extreme be faster than using existing CPU instructions without coprocessor?

Intel typically designs quite large CPU cores: 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, out-of-order unit, etc.

“So it's small cost for Intel to add instruction-set extensions for my favorite crypto!”

CPU designer's metric:

What is best performance for a specified mix of operations within a particular CPU area?

CPU designer is much more likely to consider incorporating a **small** Keccak coprocessor.

“So we should design the *smallest* Keccak circuit?”

—Maybe, but will this extreme be faster than using existing CPU instructions without coprocessor?

Intel typically designs quite large CPU cores: 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, out-of-order unit, etc.

“So it's small cost for Intel to add instruction-set extension for my favorite crypto!”

CPU designer's metric:

What is best performance for a specified mix of operations within a particular CPU area?

CPU designer is much more likely to consider incorporating a **small** Keccak coprocessor.

“So we should design the *smallest* Keccak circuit?”

—Maybe, but will this extreme be faster than using existing CPU instructions without coprocessor?

Intel typically designs quite large CPU cores: 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, out-of-order unit, etc.

“So it's small cost for Intel to add instruction-set extension for my favorite crypto!”

—Yes, but even smaller benefit for Intel's mix of operations.

designer's metric:

best performance

specified mix of operations

particular CPU area?

designer is much more likely

to incorporate a

Keccak coprocessor.

Should design the

Keccak circuit?"

Yes, but will this extreme

be better than using existing CPU

operations without coprocessor?

Intel typically designs quite large CPU cores:

32KB L1 data cache,

32KB L1 instruction cache,

several fast multipliers,

many different instructions,

out-of-order unit, etc.

"So it's small cost for Intel to add instruction-set extension for my favorite crypto!"

—Yes, but even smaller benefit for Intel's mix of operations.

Intel did

for 1 round

How many

in an AE

Can be

8: small

4: even

... 1: p

compare

and using

metric:  
 performance  
 mix of operations  
 CPU area?  
 much more likely  
 incorporating a  
 processor.  
 design the  
 circuit?"  
 this extreme  
 existing CPU  
 coprocessor?

Intel typically designs quite large CPU cores:  
 32KB L1 data cache,  
 32KB L1 instruction cache,  
 several fast multipliers,  
 many different instructions,  
 out-of-order unit, etc.

“So it’s small cost for Intel to add instruction-set extension for my favorite crypto!”

—Yes, but even smaller benefit for Intel’s mix of operations.

Intel did add instructions for 1 round of AES.  
 How many parallel in an AES-round?  
 Can be 16: big; fast  
 8: smaller but slow  
 4: even smaller but  
 ... 1: probably not  
 compared to skipping  
 and using other CPU

Intel typically designs quite large CPU cores:  
 32KB L1 data cache,  
 32KB L1 instruction cache,  
 several fast multipliers,  
 many different instructions,  
 out-of-order unit, etc.

“So it’s small cost for Intel to add instruction-set extension for my favorite crypto!”

—Yes, but even smaller benefit for Intel’s mix of operations.

Intel did add instruction for 1 round of AES.

How many parallel S-boxes in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile compared to skipping coprocessor and using other CPU instructions.

Intel typically designs quite large CPU cores: 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, out-of-order unit, etc.

“So it’s small cost for Intel to add instruction-set extension for my favorite crypto!”

—Yes, but even smaller benefit for Intel’s mix of operations.

Intel did add instruction for 1 round of AES.

How many parallel S-boxes are in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile compared to skipping coprocessor and using other CPU instructions.

Intel typically designs quite large CPU cores: 32KB L1 data cache, 32KB L1 instruction cache, several fast multipliers, many different instructions, out-of-order unit, etc.

“So it’s small cost for Intel to add instruction-set extension for my favorite crypto!”

—Yes, but even smaller benefit for Intel’s mix of operations.

Intel did add instruction for 1 round of AES.

How many parallel S-boxes are in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile compared to skipping coprocessor and using other CPU instructions.

An instruction for 4 rounds of SHA-256 is in a few Intel CPUs.

ically designs  
 ge CPU cores:  
 1 data cache,  
 1 instruction cache,  
 fast multipliers,  
 fferent instructions,  
 rder unit, etc.

small cost for Intel  
 nstruction-set extension  
 favorite crypto!”

out even smaller benefit  
 's mix of operations.

Intel did add instruction  
 for 1 round of AES.

How many parallel S-boxes are  
 in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile  
 compared to skipping coprocessor  
 and using other CPU instructions.

An instruction for 4 rounds of  
 SHA-256 is in a few Intel CPUs.

Lightwei

Frequent

where  $X$

- Keccak

- any se

- a secu

“Resourc

need the

Intel did add instruction  
for 1 round of AES.

How many parallel S-boxes are  
in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile  
compared to skipping coprocessor  
and using other CPU instructions.

An instruction for 4 rounds of  
SHA-256 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in  
where  $X$  might be

- Keccak;
  - any secure hash;
  - a secure cipher;
- “Resource-constrained  
need the smallest

Intel did add instruction  
for 1 round of AES.

How many parallel S-boxes are  
in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile  
compared to skipping coprocessor  
and using other CPU instructions.

An instruction for 4 rounds of  
SHA-256 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in literature,  
where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . . :

“Resource-constrained IoT devices  
need the smallest circuit for

Intel did add instruction for 1 round of AES.

How many parallel S-boxes are in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile compared to skipping coprocessor and using other CPU instructions.

An instruction for 4 rounds of SHA-256 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in literature, where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . :

“Resource-constrained IoT devices need the smallest circuit for  $X$ .”

Intel did add instruction for 1 round of AES.

How many parallel S-boxes are in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile compared to skipping coprocessor and using other CPU instructions.

An instruction for 4 rounds of SHA-256 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in literature, where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . :

“Resource-constrained IoT devices need the smallest circuit for  $X$ .”

—Even if speed is acceptable, who will use smallest  $X$  circuit?

Intel did add instruction  
for 1 round of AES.

How many parallel S-boxes are  
in an AES-round coprocessor?

Can be 16: big; fast.

8: smaller but slower.

4: even smaller but slower.

... 1: probably not worthwhile  
compared to skipping coprocessor  
and using other CPU instructions.

An instruction for 4 rounds of  
SHA-256 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in literature,  
where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . :

“Resource-constrained IoT devices  
need the smallest circuit for  $X$ .”

—Even if speed is acceptable,  
who will use smallest  $X$  circuit?

Why should minimum area for  $X$   
give minimum area for IoT+ $X$ ?

add instruction  
 und of AES.

ny parallel S-boxes are  
 ES-round coprocessor?

16: big; fast.

er but slower.

smaller but slower.

robably not worthwhile

ed to skipping coprocessor  
 g other CPU instructions.

uction for 4 rounds of

6 is in a few Intel CPUs.

## Lightweight crypto

Frequent claim in literature,  
 where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . .:

“Resource-constrained IoT devices  
 need the smallest circuit for  $X$ .”

—Even if speed is acceptable,  
 who will use smallest  $X$  circuit?

Why should minimum area for  $X$   
 give minimum area for IoT+ $X$ ?

## An idea

Consider  
 public ke

receives  
 under th  
 verifies t

e.g. an S

Painful h  
 all client  
 to supp

since old

## Lightweight crypto

Frequent claim in literature,  
where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . . :

“Resource-constrained IoT devices  
need the smallest circuit for  $X$ .”

—Even if speed is acceptable,  
who will use smallest  $X$  circuit?

Why should minimum area for  $X$   
give minimum area for IoT+ $X$ ?

## An idea from Adam

Consider a device  
public keys from the  
receives data supplied  
under these public  
verifies these signatures  
e.g. an SSL client.

Painful historical example  
all clients needed updates  
to support new hardware  
since old functions

## Lightweight crypto

Frequent claim in literature,  
where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . . :

“Resource-constrained IoT devices  
need the smallest circuit for  $X$ .”

—Even if speed is acceptable,  
who will use smallest  $X$  circuit?

Why should minimum area for  $X$   
give minimum area for IoT+ $X$ ?

## An idea from Adam Langley

Consider a device that receives  
public keys from trusted sources,  
receives data supposedly signed  
under these public keys;  
verifies these signatures.

e.g. an SSL client.

Painful historical event:  
all clients needed upgrades  
to support new hash functions  
since old functions were broken.

## Lightweight crypto

Frequent claim in literature,  
where  $X$  might be

- Keccak;
- any secure hash;
- a secure cipher; . . . :

“Resource-constrained IoT devices  
need the smallest circuit for  $X$ .”

—Even if speed is acceptable,  
who will use smallest  $X$  circuit?

Why should minimum area for  $X$   
give minimum area for IoT+ $X$ ?

## An idea from Adam Langley

Consider a device that receives  
public keys from trusted sources;  
receives data supposedly signed  
under these public keys;  
verifies these signatures.

e.g. an SSL client.

Painful historical event:  
all clients needed upgrades  
to support new hash functions  
since old functions were broken.

light crypto

claim in literature,

might be

k;

cure hash;

re cipher; . . . :

ce-constrained IoT devices

the smallest circuit for  $X$ ."

if speed is acceptable,

use smallest  $X$  circuit?

ould minimum area for  $X$

imum area for IoT+ $X$ ?

## An idea from Adam Langley

Consider a device that receives public keys from trusted sources; receives data supposedly signed under these public keys; verifies these signatures.

e.g. an SSL client.

Painful historical event:

all clients needed upgrades

to support new hash functions

since old functions were broken.

A public

signature

in a limi

Langley'

Replace

a full pro

Then ca

(or upgr

signature

keys, wit

## An idea from Adam Langley

Consider a device that receives public keys from trusted sources; receives data supposedly signed under these public keys; verifies these signatures.

e.g. an SSL client.

Painful historical event:  
all clients needed upgrades to support new hash functions since old functions were broken.

A public key is a signature-verification in a limited language.

Langley's idea:  
Replace this language with a full programming language. Then can upgrade (or upgrade to post-signatures!) by changing keys, with no change.

## An idea from Adam Langley

Consider a device that receives public keys from trusted sources; receives data supposedly signed under these public keys; verifies these signatures.

e.g. an SSL client.

Painful historical event:  
all clients needed upgrades to support new hash functions since old functions were broken.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash functions (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

## An idea from Adam Langley

Consider a device that receives public keys from trusted sources; receives data supposedly signed under these public keys; verifies these signatures.

e.g. an SSL client.

Painful historical event:

all clients needed upgrades to support new hash functions since old functions were broken.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

## An idea from Adam Langley

Consider a device that receives public keys from trusted sources; receives data supposedly signed under these public keys; verifies these signatures.

e.g. an SSL client.

Painful historical event:

all clients needed upgrades to support new hash functions since old functions were broken.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

Same for public-key encryption systems: public key is program.

from Adam Langley

... a device that receives  
... keys from trusted sources;  
... data supposedly signed  
... these public keys;  
... these signatures.

SSL client.

historical event:

... needed upgrades

... new hash functions

... functions were broken.

A public key is a  
signature-verification program  
in a limited language.

Langley's idea:

Replace this language with  
a full programming language.

Then can upgrade hash function  
(or upgrade to post-quantum  
signatures!) by changing public  
keys, with no changes to clients.

Same for public-key encryption  
systems: public key is program.

Say verification  
is a chip  
How small

Have to  
size of a  
size of a

Tom Langley

that receives  
trusted sources;

allegedly signed

keys;

signatures.

event:

upgrades

hash functions

as were broken.

A public key is a  
signature-verification program  
in a limited language.

Langley's idea:

Replace this language with  
a full programming language.

Then can upgrade hash function  
(or upgrade to post-quantum  
signatures!) by changing public  
keys, with no changes to clients.

Same for public-key encryption  
systems: public key is program.

Say verification de  
is a chip of area  $A$

How small can pub

Have to consider,

size of a SHA-256

size of a Keccak p

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

Same for public-key encryption systems: public key is program.

Say verification device is a chip of area  $A$ .

How small can public keys be?

Have to consider, e.g.,

size of a SHA-256 program,

size of a Keccak program, etc.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

Same for public-key encryption systems: public key is program.

Say verification device is a chip of area  $A$ .

How small can public keys be?

Have to consider, e.g.,  
size of a SHA-256 program,  
size of a Keccak program, etc.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

Same for public-key encryption systems: public key is program.

Say verification device is a chip of area  $A$ .

How small can public keys be?

Have to consider, e.g., size of a SHA-256 program, size of a Keccak program, etc.

Similar question to optimizing total size of a CPU with a SHA-256 instruction, a Keccak instruction, etc.

A public key is a signature-verification program in a limited language.

Langley's idea:

Replace this language with a full programming language.

Then can upgrade hash function (or upgrade to post-quantum signatures!) by changing public keys, with no changes to clients.

Same for public-key encryption systems: public key is program.

Say verification device is a chip of area  $A$ .

How small can public keys be?

Have to consider, e.g., size of a SHA-256 program, size of a Keccak program, etc.

Similar question to optimizing total size of a CPU with a SHA-256 instruction, a Keccak instruction, etc.

Not the usual code-size question. Change the language!